arXiv:1702.04073v1 [math.CO] 14 Feb 2017

Kneser graphs are like Swiss cheese

Ehud Friedgut* Oded Regev †

Abstract

We prove that for a large family of product graphs, and for Kneser graphs $K(n, \alpha n)$ with fixed $\alpha < 1/2$, the following holds. Any set of vertices that spans a small proportion of the edges in the graph can be made independent by removing a small proportion of the vertices of the graph. This allows us to strengthen the results of [DFR08] and [DF09], and show that any independent set in these graphs is almost contained in an independent set which depends on few coordinates. Our proof is inspired by, and follows some of the main ideas of Fox's proof of the graph removal lemma [Fox11].

1 Introduction

The celebrated triangle removal lemma of Ruzsa and Szemerédi [RS78] has a deceptively simple formulation, yet is in fact a deep structural result. It is known to imply Roth's theorem on three term arithmetic progressions, and its generalizations to hypergraphs imply Szemerédi's theorem on arithmetic progressions. The statement is as follows.

Theorem (Ruzsa-Szemerédi)**.** *For every $\varepsilon > 0$ there exists a $\delta > 0$ such that if a graph on n vertices spans less than δn^3 triangles, it can be made triangle free by removing εn^2 edges.*

This statement can be formulated as a statement regarding a 3-uniform hypergraph whose vertices are the edges of the complete graph on n vertices, and whose edges are the triangles. The statement then says that every set of vertices that spans few edges can be made independent by removing a small number of vertices.

In this paper we take this statement "one level down" to graphs. We show that a large family of graphs has an "edge removal phenomenon," i.e., any set of vertices that spans $o(|E|)$ edges can be made independent by removing $o(|V|)$ vertices. In particular, we would prove this for product graphs (see Theorem 3.1 for the precise statement), a good example of which is K_3^n, in which the set of vertices is $\{0,1,2\}^n$, and two vertices span an edge if they differ in all coordinates. (Note that this is very different from the Hamming graph, where neighbors differ in precisely one coordinate.) We will also prove this for Kneser graphs $K(n,k)$ (see Theorem 7.2), where the

*Faculty of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot 76100, Israel. Research supported in part by I.S.F. grant 1168/15 and by Minerva grant 712023.

†Courant Institute of Mathematical Sciences, New York University. This material is based upon work supported by the National Science Foundation under Grant No. CCF-1320188. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

vertices are the subsets of size k of $[n]$ for some $0 < k < n/2$, and two vertices span an edge if they are disjoint. Our proof works as long as the ratio k/n is bounded away from zero.

So, as in Swiss cheese, although these graphs have plenty of very large holes (independent sets), there are no "uniformly sparse sets": any sparse set is nothing else than a part of one of the big independent sets, with a small perturbation.

Juntas. The context in which we encountered this problem was in trying to characterize independent sets in graph products. In K_3^n, for instance, the largest independent sets are those obtained by fixing a single coordinate ("dictatorships"). Other examples of independent sets that depend on few coordinates (so called "juntas") are, say, all vertices that have at least two "0" entries in their first three coordinates. Similarly, in the Kneser graph, the largest independent sets are determined by a single "coordinate", namely, all sets containing a specific element (this is part of the Erdős-Ko-Rado theorem). And, as in K_3^n, there exist junta-like independent sets capturing a constant proportion of the vertices, which are determined by few elements.

In [DFR08] and [DF09] it is proven that *any* independent set in these graphs can essentially be captured by a set depending on few coordinates. E.g., for every $\varepsilon > 0$ there exist a $\delta > 0$ and a positive integer j such that for any independent (or very sparse) set $U \subset V(K_3^n)$ there exists a set of coordinates $J \subset [n]$ with $|J| \leq j$ and a set $T \subset \{0,1,2\}^J$ such that all but $\varepsilon 3^n$ of the vertices in U have their J-coordinates in T. Furthermore, this set T "explains" why U is independent, because T itself is extremely sparse in the graph K_3^J. A similar statement was proven there for a large class of product graphs, and for the Kneser graphs.

However, there was a fly in this ointment (or a thorn in the sheep-tail-fat, as we say): we conjectured that there must exist a set T as above that is not only sparse in the product graph of dimension j, but actually independent, thus providing a *complete* explanation for the sparseness of U; i.e., we conjectured that any independent (or very sparse) set is almost completely contained in a set depending on few coordinates that is truly independent (as opposed to merely sparse). In this paper we manage to settle this issue and prove the conjecture. The key to this is applying the main theorem in this paper, Theorem 3.1, to the set T, which belongs to the j'th power of the base graph. We show that the edge removal phenomenon holds in the product graph, thus T can be slightly altered to produce a *truly independent set* (as opposed to a sparse one).

Related work. The structure of our proof is very closely modeled on Fox's proof of the graph removal lemma, [Fox11], where he improved the longstanding bound on the dependence of the constants in the lemma (most famously, in the triangle removal lemma). One of the differences is that we have, to use Fox's terminology, a shattering lemma that is special to this setting, and is nothing else than a (minor) generalization of the main result of [DFR08] (Theorem 2.1). It states that whenever two large sets of vertices span few edges between them, it is because there is a small set of coordinates that these two sets are strongly correlated with. This is a consequence of a central theorem in [DMR09], which in turn, relies heavily on the invariance principle of [MOO10].

Our main result is also reminiscent of, and related to, removal lemmas in groups, see [HST16], [KSV13], and for systems of equations over finite fields, [KSV12]. Theorem 7.2, which deals with Kneser graphs, is closely related to the work [DT16], where a removal lemma is proven for the special case of sets that are close in size to a maximal independent set in the Kneser graph $K(n,k)$, for any value of k.

Structure of the paper. In Section 2 we present some preliminary definitions. In Section 3 we state our main theorem and sketch the proof. In Section 4 we reduce to a case of a "matching like" function, which will eventually lead to a substantial improvement in the resulting bounds. In Section 5 we show that any non-negative function has a good approximation (in a specific sense) by a function that depends on few coordinates. In Section 6 we complete the proof of our main theorem. In Section 7 we state and prove the Swiss Cheese Theorem regarding Kneser graphs, showing that they too exhibit an edge removal phenomenon. Finally, in the appendix we explain how to extend the main theorem from [DFR08] to the form that we use in this paper.

2 Preliminaries

In the rest of the paper we fix a set V with a reversible, irreducible, aperiodic Markov chain on it given by a matrix A. All functions and constants we encounter from now on may depend on V and A. Let μ denote the unique stationary measure of A on V. We will be working with V^n, $A^{\otimes n}$ and $\mu^{\otimes n}$, often just writing A and μ as shorthand to avoid cumbersome notation. Whenever we take expectation of a function on V^n it is according to μ, and we use μ also to define the standard inner product between functions on V^n. We think of the ground set V^n as the vertices of a (product) graph, and of the non-zero-probability transitions as edges. The weight of a (directed) edge (x, y) is

$$w_{(x,y)} := \langle \mathbb{1}_x, A\mathbb{1}_y \rangle .$$

This is the asymptotic probability of a step in the random walk governed by A to traverse (x, y). Equivalently it is the probability of $(u, w) = (x, y)$, where u is chosen by the stationary distribution, and w is chosen from u's neighbors according to the probabilities dictated by the transition matrix. We will talk about "the weight of the edges spanned" by a set U, or a function f, simply meaning $\langle \mathbb{1}_U, A\mathbb{1}_U \rangle$, or $\langle f, Af \rangle$. Consequently, a set $U \subset V^n$ is called *independent* if $\langle \mathbb{1}_U, A\mathbb{1}_U \rangle = 0$. We will say that a function $g : V^n \to [0, 1]$ is ε-far from independent if for every independent set U we have $\mathbb{E}[\mathbb{1}_{\overline{U}} \cdot g] > \varepsilon$.

Capturing sparse sets using juntas. A crucial ingredient in our proof is the following variant on the main result from [DFR08]. In the appendix we explain how it follows from previous work.

Theorem 2.1. *For all $0 < \varepsilon \leq 1/2$ there exist $\delta > 0, j \geq 1$, such that the following holds. For all $n \geq 1$ and $f_1, f_2 : V^n \to [0, 1]$ such that $\langle f_1, Af_2 \rangle \leq \delta$, there exist $J \subseteq [n]$, $|J| \leq j$, and $T_1, T_2 \subseteq V^J$ such that*

$$\mathbb{E}_{x \in V^J}[\mathbb{1}_{\overline{T_1}}(x)\, \mathbb{E}[f_1(x, \cdot)]] \leq \varepsilon, \qquad \mathbb{E}_{x \in V^J}[\mathbb{1}_{\overline{T_2}}(x)\, \mathbb{E}[f_2(x, \cdot)]] \leq \varepsilon, \tag{1}$$

and $\langle \mathbb{1}_{T_1}, A\mathbb{1}_{T_2} \rangle \leq \varepsilon$. Moreover, one can take $\delta = \delta_1(\varepsilon) := \varepsilon^c$ and $j = j_1(\varepsilon) := \varepsilon^{-c}$ where $c > 0$ is a constant depending only on A.

3 Main Theorem

Theorem 3.1. *For all $\varepsilon > 0$ there exists $\delta > 0$ such that the following holds. If $g : V^n \to [0, 1]$ is ε-far from independent then $\langle g, Ag \rangle > \delta$. Moreover, we can take $\delta = \delta_2(\varepsilon) := 1/\mathrm{Tower}(O(\log(1/\varepsilon)))$.*

This easily implies the desired strengthening of [DFR08] which was the main motivation for this work.

Corollary 3.2. *For all $\varepsilon > 0$ there exist a $\delta > 0$ and a positive integer j such that the following holds. Let $g : V^n \to [0,1]$, with $\langle g, Ag \rangle < \delta$. Then there exists $J \subset [n]$, with $|J| \leq j$, and $T \subseteq V^J$ such that*

1. $\langle \mathbb{1}_T, A\mathbb{1}_T \rangle = 0$, *and*
2. $\mathbb{E}_{x \in V^J}[\mathbb{1}_{\overline{T}}(x)\, \mathbb{E}[g(x,\cdot)]] \leq \varepsilon$.

Moreover, we can take $\delta = \delta_1(\delta_2(\varepsilon/2))$ and $j = j_1(\delta_2(\varepsilon/2))$.

Proof. Invoke Theorem 2.1 with $f_1 = f_2 = g$ and ε taken to be $\delta_2(\varepsilon/2)$ to produce J, with $|J| \leq j$, and $T' \subset V^J$ satisfying that

$$\mathbb{E}_{x \in V^J}[\mathbb{1}_{\overline{T'}}(x)\, \mathbb{E}[g(x,\cdot)]] \leq \delta_2(\varepsilon/2) \leq \varepsilon/2,$$

and $\langle \mathbb{1}_{T'}, A\mathbb{1}_{T'} \rangle \leq \delta_2(\varepsilon/2)$. Then invoke Theorem 3.1 on $\mathbb{1}_{T'}$ to find an independent set $T \subset T'$ with

$$\mathbb{E}[\mathbb{1}_{T'} - \mathbb{1}_T] \leq \varepsilon/2\ .$$

It follows that T satisfies conditions (1) and (2), as required. □

3.1 Sketch of proof of the main theorem

We begin the sketch for the case of sparse sets (i.e., for g having range $\{0,1\}$), as that captures all the main ideas. We will briefly mention the extension to functions at the end of this subsection. Given a set $U \subset V^n$ which is ε-far from being independent we wish to show that it spans many edges, i.e., that $\langle \mathbb{1}_U, A\mathbb{1}_U \rangle$ is large.

Notice that any set $I \subset [n]$ naturally defines a partition of V^n into $|V|^{|I|}$ parts according to the coordinates in I. We can then study how $\mathbb{1}_U$ behaves on these parts. If it is constant (zero or one) on each part, then I perfectly captures U. Letting W be a random variable whose value is the conditional expectation of $\mathbb{1}_U$ on a random part of the partition, and letting $H := H(U,I)$ be the expectation of $W \log(W)$, then H is a good indication of how well I captures U. When $I = \emptyset$ then $H = \alpha \log(\alpha)$, where α is the measure of U. Moreover, $H = 0$ if and only if $\mathbb{1}_U$ is completely determined by the coordinates in I. Furthermore, H is always non-positive, and is monotone increasing with respect to refining the partition induced by I by adding further coordinates. For a positive integer r, an r-refinement of a partition induced by I as above, is one which involves adding r new coordinates per every part of the partition, thus adding up to $r \cdot |V|^{|I|}$ new coordinates.

Given a set U, beginning with the trivial partition of V^n (corresponding to $I_0 = \emptyset$) we will iteratively apply r-refinements, producing $I_0 \subset I_1 \subset \cdots$ attempting to substantially increase $H(U, I_i)$ in each step, and stop when this is no longer possible. On the one hand, since $H \leq 0$, the number of steps, and hence the total number of coordinates involved in the final partition is bounded from above by some constant k. On the other hand, we will show that if U is ε-far from an independent set, and $\langle \mathbb{1}_U, A\mathbb{1}_U \rangle = \delta < \delta(\varepsilon, k)$ then for any partition coming from at most k coordinates there is an r-partition that further increases H substantially. So if δ is too small this yields a contradiction. In the proof we will sketch the exact dependence between all parameters involved (including r and k).

The crux of the proof, then, is how one can utilize the fact that U is sparse (i.e., that δ is small), U is ε-far from being independent, and $|I|$ is not too large, in order to show the existence of an r-refinement which substantially increases $H(U,I)$. Our engine for this is (a slight variation on)

the result from [DFR08], where the fuel of this engine is the invariance principle of [MOO10] (as applied in [DMR09]). Whenever two parts of the partition, say X and Y, span few edges between them, our engine will produce a refinement of the partition, according to a bounded number of new coordinates, such that on at least one of the two parts, say X, the resulting increase in H will be proportional to the measure of $U \cap X$. This approach is sufficient to prove our main theorem, with a Tower$(O(1/\varepsilon))$-type dependence between ε and δ.

However, using a key idea from Fox's improvement to the bounds in the graph removal lemma, [Fox11], we can cut this dependence down to Tower$(O(\log(1/\varepsilon)))$. This involves replacing the set U with a subset U', the support of a maximal matching contained in it. The advantage of U' is that no independent set contained in it captures more than half of its mass. As the proof will show, this helps avoid slowly whittling away at U, and speeds up the relative increase in H in every step.

One last element in the proof of our main theorem: the use of functions with range $[0,1]$ instead of sets. In the final part of the paper, when applying our result to sets in Kneser graphs, we will "extrapolate" these sets to $[0,1]$-valued functions on $\{0,1\}^n$. Therefore it makes sense for us to make a few small adaptations in our presentation, replacing sets (which may be thought of as functions with range $\{0,1\}$) with functions with range $[0,1]$. It turns out that this natural variant is not much harder to treat than the original one.

4 Reducing to matching-like functions

Here we observe that it suffices to prove Theorem 3.1 in the special case that f is a *matching-like function*, which is a function not having too much weight on any independent set. The use of this innocuous condition, which replaces the condition of being far from independent, leads to a substantial improvement.

Definition 4.1. *Let $f : V^n \to [0,1]$. We say that f is* matching-like *if for any independent set $W \subseteq V^n$,*

$$\mathbb{E}_x[\mathbb{1}_W(x)f(x)] \leq \mathbb{E}[f]/2.$$

The terminology comes from the observation that when μ is the uniform measure, the indicator function of the vertices touched by any matching is "matching-like," since no independent set can contain both ends of an edge.

The following claim shows that for any function $g : V^n \to [0,1]$ we can find a matching-like function f such that $g \geq f$ pointwise, and such that the set of vertices x such that $f(x) < g(x)$ is an independent set. We note that when the measure μ is the uniform measure, and g is the indicator of a set U, then we can simply take any maximal matching inside U and let f be the indicator of its vertices. The three properties below are then easy to verify.

Claim 4.2. *For any $g : V^n \to [0,1]$ there exists a function $f : V^n \to [0,1]$ satisfying that*

1. *$f(x) \leq g(x)$ for all x.*
2. *f is matching-like, and*
3. *$\{x : f(x) < g(x)\}$ is an independent set.*

Proof. Note that the set of all functions f that fulfill conditions (1) and (2) is non-empty (as it contains the identically 0 function), and compact. Hence there exists some f in this set which maximizes $\sum f(x)$. This f must also fulfill condition (3), as otherwise we would have an edge $\{a, b\}$ with $f(a) < g(a)$ and $f(b) < g(b)$. If such an edge were to exist, there would exist some small positive constant γ such that we could add $\gamma/\mu(a)$ to $f(a)$, and $\gamma/\mu(b)$ to $f(b)$, yielding a new function f' that still fulfills condition (1) and (3). This function will also fulfill condition (2) since its expectation is greater than that of f by 2γ, but its weight on any independent set is greater by at most γ (because no independent set contains both a and b.) So the existence of f' contradicts the maximality of f. □

We now state a theorem quite similar to our main one, Theorem 3.1, for the special case of matching-like functions, and then, using Claim 4.2, we can easily deduce Theorem 3.1.

Theorem 4.3. *For all $\varepsilon > 0$ there exists a $\delta > 0$ such that any matching-like $f : V^n \to [0,1]$ with $\mathbb{E}[f] \geq \varepsilon$ satisfies $\langle f, Af \rangle > \delta$. Moreover, we can take $\delta = 1/\mathrm{Tower}(O(\log(1/\varepsilon)))$.*

This now easily implies our main theorem.

Proof of Theorem 3.1. Given g we invoke Claim 4.2 to produce an appropriate matching-like f. Note that since the set $U := \{x : f(x) < g(x)\}$ is independent we have that

$$\mathbb{E}[f] \geq \mathbb{E}[\mathbb{1}_{\overline{U}} \cdot f] = \mathbb{E}[\mathbb{1}_{\overline{U}} \cdot g] > \varepsilon.$$

So, by Theorem 4.3 we have $\langle f, Af \rangle > \delta$. Since $g \geq f$ pointwise it follows that $\langle g, Ag \rangle > \delta$ as required. □

5 The potential argument

Given a function $f : V^n \to [0,1]$, and a set of coordinates $I \subseteq [n]$ we wish to study how well $f(x)$ is predicted by the coordinates of x indexed by I. If f depends only on the coordinates in I, then we think of f as being perfectly correlated with I, in which case f is constant on every part of the partition of V^n induced by I. Otherwise, we can improve this correlation by refining the partition according to additional coordinates. In Definition 5.1 we set a (partially arbitrary) benchmark for how much this refinement improves the correlation, and define any refinement that succeeds as "substantially improving the correlation". If for every x in V^I we partition $\{x\} \times V^{[n]\setminus I}$ according to r additional coordinates we call this an "r-refinement". Our main goal in this section is the proof of Lemma 5.2 that states, roughly, that for any f and r there exists a (not-too-large) set of coordinates I whose correlation with f cannot be substantially improved by r-refinement. Definition 5.1 and the statement of Lemma 5.2 are the only parts of this section used in the rest of the paper.

Definition 5.1. *Consider a function $f : V^n \to [0,1]$ and let α denote $\mathbb{E}[f]$. For $r \geq 1$ and for a set of coordinates $I \subseteq [n]$, we say that the correlation of I with f can be* substantially improved by r-refinement *if there exists a subset $S \subseteq V^I$ for which the following holds:*

1. $\mathbb{E}_{x \in V^I}[\mathbb{1}_S(x)\, \mathbb{E}[f(x, \cdot)]] \geq \alpha/2$, *and*
2. *For each $x \in S$ there exists $J_x \subseteq [n] \setminus I$ of cardinality at most r and $T_x \subseteq V^{J_x}$ satisfying*

(a) $\Pr[y \in T_x] \leq 3/4$, *and*
(b) $\mathbb{E}_{y \notin T_x}[\mathbb{E}[f(x,y,\cdot)]] \leq \alpha/8$.

Just to ensure the notation in (*b*) is clear: $\mathbb{E}_{y \notin T_x}[h(y)] := \frac{\sum_{y \notin T_x} \mu_p(y) h(y)}{\Pr[y \notin T_x]}$.

As mentioned, a set of coordinates $I \subset [n]$ is perfectly correlated with f if for every $x \in V^I$ it holds that $f(x,\cdot)$ is constant. Substantially improving the correlation of I with f by refinement, means that for a portion of inputs $x \in V^I$, which are responsible for at least half of the expectation of f, it holds that by an appropriate choice of J_x, (at most r additional coordinates from $[n] \setminus I$), partitioning $\{x\} \times V^{[n] \setminus I}$ according to these additional coordinates yields a partition where in many parts the conditional expectation of f drops substantially. In a sense this implies that f is closer to being constant on the parts of the refined partition of $V^{[n]}$.

Lemma 5.2. *Consider a function $f : V^n \to [0,1]$ and let α denote $\mathbb{E}[f]$. Let r be a positive integer, and define the function $\Gamma(\ell) := \ell + r|V|^\ell$. Then there exists a set $J \subseteq [n]$ of cardinality at most $k = k(\alpha, r) := \Gamma^{\circ 128 \log(1/\alpha)}(0)$ whose correlation with f cannot be substantially improved by r-refinement, where $\Gamma^{\circ t}$ denotes the composition of Γ with itself t times.*

The proof uses a potential argument. Define the function $\varphi : \mathbb{R}^{\geq 0} \to \mathbb{R}$ by $\varphi(x) = x \log x$ (and set $\varphi(0) = 0$.) This is a convex function and is non-positive on $[0,1]$. For $f : V^n \to [0,1]$ and $I \subseteq [n]$ define the *entropy* of f with respect to I as

$$H(f, I) := \mathbb{E}_{x \in V^I}[\varphi(\mathbb{E}[f(x,\cdot)])].$$

Since φ is convex, we get from Jensen's inequality that $H(f,I)$ is monotone in I, i.e., $H(f,I) \leq H(f,J)$ whenever $I \subseteq J$. Moreover, $H(f,I) \leq 0$ for all I.

Claim 5.3. *Let $f : V^n \to [0,1]$ be a function and $\alpha = \mathbb{E}[f]$. Let $I \subseteq [n]$. If the correlation of I with f is substantially improved by r-refinement for some $r \geq 1$ then*

$$H(f, J) \geq H(f, I) + \alpha/128,$$

where $J := I \cup \bigcup_{x \in S} J_x$, and S, J_x are as in Definition 5.1.

We first observe that this claim implies Lemma 5.2.

Proof of Lemma 5.2. Notice that by definition $H(f, \emptyset) = \alpha \log \alpha$. Hence, starting with $I_0 = \emptyset$, we repeatedly apply Claim 5.3, and create an r-refinement, each time obtaining a new set I_{t+1} with $H(f, I_{t+1}) \geq H(f, I_t) + \alpha/128$, and $|I_{t+1}| \leq \Gamma(|I_t|)$. We continue as long as the correlation of I_{t+1} with f can be substantially improved by r-refinement. Since $H(f,I) \leq 0$ for all I, this process must terminate after at most $128 \log(1/\alpha)$ steps, implying the lemma. □

Proof of Claim 5.3. Let S, J_x, T_x be as in Definition 5.1. Define $S' \subseteq S$ as the set of $x \in S$ satisfying $\mathbb{E}[f(x,\cdot)] \geq \alpha/4$. By Item 1 of Definition 5.1,

$$\mathbb{E}_{x \in V^I}[\mathbb{1}_{S'}(x)\, \mathbb{E}[f(x,\cdot)]] \geq \mathbb{E}_{x \in V^I}[\mathbb{1}_S(x)\, \mathbb{E}[f(x,\cdot)]] - \alpha/4 \geq \alpha/4. \tag{2}$$

Now,

$$\begin{aligned} H(f,J) &= \mathbb{E}_{x \in V^I}\left[H\Big(f(x,\cdot), \bigcup_{s \in S} J_s\Big)\right] \\ &= \mathbb{E}_{x \in V^I}\left[\mathbb{1}_{S'}(x) H\Big(f(x,\cdot), \bigcup_{s \in S} J_s\Big)\right] + \mathbb{E}_{x \in V^I}\left[\mathbb{1}_{\overline{S'}}(x) H\Big(f(x,\cdot), \bigcup_{s \in S} J_s\Big)\right]. \end{aligned} \tag{3}$$

By monotonicity, the second term in (3) is at least

$$\mathbb{E}_{x\in V^I}[\mathbb{1}_{\overline{S'}}(x)\varphi(\mathbb{E}[f(x,\cdot)])]. \tag{4}$$

For analyzing the first term in (3), fix any $x \in S'$. Then using monotonicity,

$$\begin{aligned}
H\Big(f(x,\cdot), \bigcup_{s\in S} J_s\Big) &\geq H(f(x,\cdot), J_x) \\
&= \mathbb{E}_{y\in V^{I_x}}[\varphi(\mathbb{E}[f(x,y,\cdot)])] \\
&= \Pr[y \notin T_x]\,\mathbb{E}_{y\notin T_x}[\varphi(\mathbb{E}[f(x,y,\cdot)])] + \Pr[y \in T_x]\,\mathbb{E}_{y\in T_x}[\varphi(\mathbb{E}[f(x,y,\cdot)])] \\
&\geq \Pr[y \notin T_x]\varphi(\mathbb{E}_{y\notin T_x}[\mathbb{E}[f(x,y,\cdot)]]) + \Pr[y \in T_x]\varphi(\mathbb{E}_{y\in T_x}[\mathbb{E}[f(x,y,\cdot)]]) \\
&= \lambda\varphi(u) + (1-\lambda)\varphi(v) \\
&\geq \varphi(w) + w/32, \qquad (5)
\end{aligned}$$

where we define $\lambda := \Pr[y \notin T_x]$, $u := \mathbb{E}_{y\notin T_x}[\mathbb{E}[f(x,y,\cdot)]]$, $v := \mathbb{E}_{y\in T_x}[\mathbb{E}[f(x,y,\cdot)]]$, and $w := \lambda u + (1-\lambda)v = \mathbb{E}[f(x,\cdot)]$. The last inequality follows from Claim 5.4 below, noting that by Definition 5.1, $\lambda \geq 1/4$ (Item 2.a), $u \leq \alpha/8$ (Item 2.b), and $w \geq \alpha/4$ by our choice of S', and therefore $u \leq w/2$ as required. Plugging this into (3), we obtain that

$$\begin{aligned}
H(f,J) &\geq \mathbb{E}_x[\mathbb{1}_{S'}(x)\varphi(\mathbb{E}[f(x,\cdot)])] + \mathbb{E}_x[\mathbb{1}_{\overline{S'}}(x)\varphi(\mathbb{E}[f(x,\cdot)])] + \frac{1}{32}\,\mathbb{E}_x[\mathbb{1}_{S'}(x)\,\mathbb{E}[f(x,\cdot)]] \\
&= H(f,I) + \frac{1}{32}\,\mathbb{E}_x[\mathbb{1}_{S'}(x)\,\mathbb{E}[f(x,\cdot)]] \\
&\geq H(f,I) + \frac{\alpha}{128},
\end{aligned}$$

where the last inequality uses (2). □

Claim 5.4. *For any $u, v > 0$ and $\lambda \in [1/4, 1]$ satisfying that $u \leq w/2$ where $w := \lambda u + (1-\lambda)v$,*

$$\lambda\varphi(u) + (1-\lambda)\varphi(v) \geq \varphi(w) + w/32\,.$$

Proof. From the definition of φ we have

$$\lambda\varphi(u) + (1-\lambda)\varphi(v) = \varphi(w) + w\big(\lambda\varphi(u/w) + (1-\lambda)\varphi(v/w)\big)\,. \tag{6}$$

Notice that $v/w = (1-\lambda u/w)/(1-\lambda) \geq 7/6$. Since φ is convex, the line segment connecting $(u/w, \varphi(u/w))$ with $(v/w, \varphi(v/w))$ lies above the line segment connecting $(1/2, \varphi(1/2))$ with $(7/6, \varphi(7/6))$ (see Figure 1), and therefore, their points of intersection with the vertical line $x = 1$ satisfy

$$\lambda\varphi(u/w) + (1-\lambda)\varphi(v/w) \geq \frac{1}{4}\varphi(1/2) + \frac{3}{4}\varphi(7/6) > 1/32\,. \tag{7}$$

Combining Eqs. (6) and (7) yields the result. □

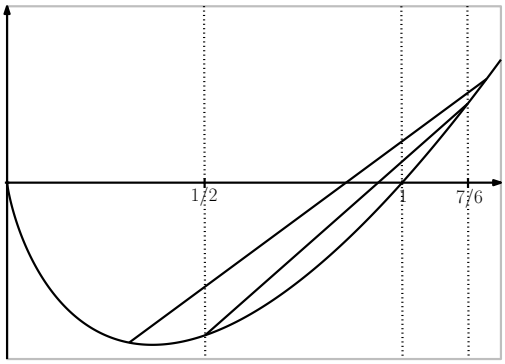

Figure 1: $\varphi(x)$ for $x \in [0, 5/4]$.

6 Improving correlation through refinements

In this section we prove Theorem 4.3, thereby completing the proof of our main theorem. It will be convenient to use the following corollary of Theorem 2.1, showing that if two functions f_1, f_2 span very few edges, then one of them must be concentrated on a junta of measure $3/4$.

Corollary 6.1. *For all $0 < \varepsilon \leq 1/2$ there exist $\delta > 0, r \geq 1$, such that the following holds. For all $n \geq 1$ and $f_1, f_2 : V^n \to [0,1]$ such that $\langle f_1, Af_2 \rangle \leq \delta$, there exist $i \in \{1,2\}$, $J \subseteq [n]$, $|J| \leq r$, and $T \subseteq V^J$ such that*

$$\mathbb{E}_{x \in V^J}[\mathbb{1}_{\overline{T}}(x)\, \mathbb{E}[f_i(x, \cdot)]] \leq \varepsilon, \tag{8}$$

and

$$\Pr[x \in T] \leq 3/4. \tag{9}$$

Moreover, one can take $\delta = \delta_1(\varepsilon) := \varepsilon^c$ and $r = r_1(\varepsilon) := \varepsilon^{-c}$ where $c > 0$ is a constant depending only on A.

Proof. Apply Theorem 2.1, yielding $|T_1|, |T_2| \leq r$ where $r = j$, and observe that $\langle \mathbb{1}_{T_1}, A\mathbb{1}_{T_2} \rangle \leq \varepsilon \leq 1/2$ implies that either $\Pr[x \in T_1] \leq 3/4$ or $\Pr[x \in T_2] \leq 3/4$. □

Let $w_{\min}$ be the minimal weight of an edge in the transition graph of A, i.e., the minimal positive value of $\langle \mathbb{1}_u, A\mathbb{1}_v \rangle$, where u and v are elements of V. Note that the weight of the minimal edge for $A^{\otimes m}$ is $w_{\min}^m$.

Lemma 6.2. *For all $\varepsilon > 0$ there exists $r \geq 1$, such that for any $k \geq 1$ there exists $\delta > 0$ such that for any matching-like $f : V^n \to [0,1]$ with $\langle f, Af \rangle \leq \delta$ and $\mathbb{E}[f] = \varepsilon$, and for any $I \subseteq [n]$ of cardinality at most k, the correlation of I with f can be substantially improved by r-refinement. Moreover, we can take $\delta = \delta_2(\varepsilon, k) := w_{\min}^k \delta_1(\varepsilon/32) = w_{\min}^k (\varepsilon/32)^c$ and $r = r_2(\varepsilon) := r_1(\varepsilon/32) = (\varepsilon/32)^{-c}$.*

Notice that r_2 depends on ε but not on k. Also note that the factor $w_{\min}^k$, which, as we will see in the proof, is incurred by the normalization when moving from V^n to V^{n-k}, is the dominant factor in our calculations (as we will need $k = \mathrm{Tower}(O(\log(1/\varepsilon)))$).

Proof. Fix a function f and a set $I \subseteq [n]$, with $|I| \leq k$ as in the statement of the lemma. We construct a set $S \subseteq V^I$ and sets J_x, T_x satisfying the requirements in Definition 5.1 as follows. Initially, we

set $S = \emptyset$. We consider all edges (including loops!) (x_1, x_2) (i.e., all pairs $(x_1, x_2) \in V^I \times V^I$ with $w_{(x_1,x_2)} > 0$), in an arbitrary order. For each edge (x_1, x_2) if either x_1 or x_2 is already in S, we continue to the next edge. Otherwise, we apply Corollary 6.1 to the functions $f(x_1, \cdot)$ and $f(x_2, \cdot)$ with ε taken to be $\varepsilon/32$, resulting in $i \in \{1, 2\}$, $J \subseteq [n] \setminus I$, with $|J| \leq r_1(\varepsilon/32)$, and $T \subseteq V^J$. We then add x_i to S, and define J_{x_i} to be J and T_{x_i} to be T. This completes the description of the construction. Notice that we are allowed to apply Corollary 6.1 above with parameter ε set to $\varepsilon/32$, since

$$\langle f(x_1, \cdot), Af(x_2, \cdot)\rangle \leq w_{\min}^{-|I|}\langle f, Af\rangle \leq w_{\min}^{-k}\delta = \delta_1(\varepsilon/32) .$$

Moreover, notice that $\overline{S}$ forms an independent set (since S must contain at least one vertex of each edge), and since f is matching-like, we have

$$\mathbb{E}_{x\in V^I}[\mathbb{1}_S(x)\,\mathbb{E}[f(x, \cdot)]] = \varepsilon - \mathbb{E}_{x\in V^I}[\mathbb{1}_{\overline{S}}(x)\,\mathbb{E}[f(x, \cdot)]] \geq \varepsilon/2.$$

Finally, for all $x \in S$ we have by (9) that $\Pr[y \in T_x] \leq 3/4$, and by (8) that

$$\mathbb{E}_{y\notin T_x}[\mathbb{E}[f(x, y, \cdot)]] = \mathbb{E}_y[\mathbb{1}_{\overline{T_x}}(y)\,\mathbb{E}[f(x, y, \cdot)]]/\Pr_y[y \notin T_x] \leq \frac{\varepsilon/32}{1/4} = \varepsilon/8.$$

We conclude that S, $\{J_x\}$, and $\{T_x\}$ satisfy the requirements in Definition 5.1, as required. □

Theorem 4.3 follows from Lemma 5.2 and Lemma 6.2.

Proof of Theorem 4.3. Given a matching-like f, with $\mathbb{E}[f] = \varepsilon$, let $r = r_2(\varepsilon)$ be as given by Lemma 6.2. Apply Lemma 5.2 with f and r to get a set J of cardinality at most $k = k(r, \varepsilon) = \text{Tower}(O(\log(1/\varepsilon)))$ such that the correlation of J with f cannot be substantially improved by r-refinement. From Lemma 6.2 it follows that $\langle f, Af\rangle > \delta$, with $\delta = \delta_3(\varepsilon, k) = (\text{Tower}(O(\log(1/\varepsilon))))^{-1}$. □

7 Kneser graphs are like Swiss cheese

In this section we prove Theorem 7.2, which extends our main theorem to the case of Kneser graphs. Fix $0 < p < 1/2$. We will consider the Markov chain on $\{0, 1\}^n$ which moves independently on each coordinate according to the transition matrix

$$\begin{pmatrix} \frac{1-2p}{1-p} & \frac{p}{1-p} \\ 1 & 0 \end{pmatrix}.$$

The stationary measure of this Markov chain is the product measure $\mu_p = (1 - p, p)^{\otimes n}$, and all transitions (x, y) have probability 0 if x and y are not disjoint. If x and y are disjoint, then the weight of the edge (x, y) is precisely $p^{|x|}p^{|y|}(1 - 2p)^{n-|x|-|y|}$. So, for two disjoint sets $x, y \subset [n]$, we define

$$\mu_{p,p}(x, y) := p^{|x|}p^{|y|}(1 - 2p)^{n-|x|-|y|}.$$

Recall the following notation. Given a set of coordinates $J \subset [n]$, and two vectors $w \in \{0, 1\}^J$, and $x \in \{0, 1\}^{[n]\setminus J}$, we will write (w, x) for the element of $\{0, 1\}^n$ formed by merging them appropriately. The following, then, is a special case of Corollary 3.2.

Theorem 7.1. *Let $0 < p < 1/2$. There exists functions $\delta_p = \delta : [0,1] \to [0,1]$ and $j_p = j : [0,1] \to \mathbb{N}$ such that the following holds. Let $g : \{0,1\}^n \to [0,1]$, and let*

$$\mathrm{Edge}(g) = \sum_{x \cap y = \emptyset} g(x)g(y)\mu_{p,p}(x,y).$$

Then for every $\varepsilon \in [0,1]$, if $\mathrm{Edge}(g) \leq \delta(\varepsilon)$ then there exists $J \subset [n]$ with $|J| \leq j(\varepsilon)$, and $T \subset \{0,1\}^J$ such that

1. *T is an intersecting family.*
2. $$\mathbb{E}_{w \in \{0,1\}^J}[\mathbb{1}_{\overline{T}}(w)\, \mathbb{E}_{x \in \{0,1\}^{[n]\setminus J}}[g(w,x)]] \leq \varepsilon,$$
   where all expectations are taken with respect to μ_p.

The main theorem of this section is very similar, except it is set on a single layer of the cube $\{0,1\}^n$, i.e., on $\binom{[n]}{k}$.

Theorem 7.2. *Let $0 < p < 1/2$ and let $\delta_p = \delta : [0,1] \to [0,1]$ and $j_p = j : [0,1] \to \mathbb{N}$ be as in Theorem 7.1. Let n and $k = pn$ be positive integers, and let $G = G(n,k)$ be the Kneser graph with*

$$V(G) = \binom{[n]}{k}, E(G) = \{\{x,y\} : x \cap y = \emptyset\}.$$

Let $f : V(G) \to [0,1]$, and let

$$\mathrm{Edge}(f) := \sum_{\{x,y\} \in E(G)} \frac{f(x)f(y)}{\binom{n}{k}\binom{n-k}{k}}.$$

Then for every $\varepsilon \in [0,1]$, and $\mathrm{Edge}(f) \leq \delta(\varepsilon)$ then if n is sufficiently large there exists $J \subset [n]$ with $|J| \leq j(\varepsilon)$, and $T \subset \{0,1\}^J$ such that

1. *T is an intersecting family, and*
2. $$\binom{n}{k}^{-1} \sum_{(x \cap J) \notin T} f(x) \leq 5\varepsilon\,.$$

We will show how to deduce Theorem 7.2 from Theorem 7.1. First, we need two lemmas, regarding moving from functions on a single layer to functions on the whole cube, and vice versa.

Lemma 7.3 (The Up Lemma)**.** *Let $k = pn$ for $0 < p < 1/2$. For $f : \binom{[n]}{k} \to [0,1]$, define $g : \{0,1\}^n \to [0,1]$ by*

$$g(x) := \begin{cases} \binom{|x|}{k}^{-1} \sum_{x' \subseteq x} f(x') & |x| \geq k \\ 0 & \text{otherwise.} \end{cases}$$

Then

$$\mathrm{Edge}(g) \leq \mathrm{Edge}(f)\,.$$

Proof. For any f and the corresponding g,

$$\begin{aligned}\mathrm{Edge}(g) &= \sum_{x\cap y=\emptyset} \mu_{p,p}(x,y)\left(\binom{|x|}{k}\binom{|y|}{k}\right)^{-1} \sum_{x'\subseteq x, y'\subseteq y} f(x')f(y') \\ &= \sum_{x'\cap y'=\emptyset} c(p,n)\frac{f(x')f(y')}{\binom{n}{k}\binom{n-k}{k}} = c(p,n)\mathrm{Edge}(f)\ ,\end{aligned}$$

where

$$c(p,n) = \sum_{x\cap y=\emptyset, x'\subseteq x, y'\subseteq y} \mu_{p,p}(x,y)\frac{\binom{n}{k}\binom{n-k}{k}}{\binom{|x|}{k}\binom{|y|}{k}}$$

is independent of the pair (x',y') and of the function f. Plugging in the case $f\equiv 1$, where $\mathrm{Edge}(f)=1$, $g(x)=\mathbf{1}_{|x|\geq k}$ gives

$$c(p,n) = \mathrm{Edge}(g) \leq 1,$$

since $g\leq 1$ □

Lemma 7.4 (The Down Lemma)**.** *Let $k=pn$, let f and g be as above, and let $J\subset[n]$, and $w\in\{0,1\}^J$. Define*

$$V_w(g) := p^{|w|}(1-p)^{|J|-|w|}\,\mathbb{E}_{x\in\{0,1\}^{[n]\setminus J}}[g(w,x)] = \sum_{x\in\{0,1\}^{[n]\setminus J}} g(w,x)\mu_p(w,x)\ .$$

Let

$$V_w(f) := \sum_x \frac{f(w,x)}{\binom{n}{k}},$$

where the sum is over $x\in\{0,1\}^{[n]\setminus J}$ of size precisely $k-|w|$. Then, for sufficiently large n,

$$V_w(f) \leq 5V_w(g)\ .$$

Proof. Throughout this proof, when summing over x, we are restricting ourselves to the case $|x|+|w|\geq k$, since other values of x contribute nothing. Observe that

$$\begin{aligned}V_w(g) &= \sum_x g(w,x)\mu_p(w,x) \\ &= \sum_x \frac{\mu_p(w,x)}{\binom{|x|+|w|}{k}} \sum_{w'\subseteq w, x'\subseteq x} f(w',x') \\ &\geq \sum_x \frac{\mu_p(w,x)}{\binom{|x|+|w|}{k}} \sum_{x'\subseteq x} f(w,x') \\ &= \sum_{x'} \frac{f(w,x')}{\binom{n}{k}}\left(\sum_{x\supseteq x'} \frac{\mu_p(w,x)\binom{n}{k}}{\binom{|x|+|w|}{k}}\right)\ . \end{aligned} \tag{10}$$

where in the last equality we reversed the order of summation. Since the first sum in Eq. (10) is precisely $V_w(f)$, it suffices to show that the second sum (which only depends on p, $|J|$, and $|w|$) is at least $1/5$. To this end observe that

$$\sum_{x\supseteq x'} \frac{\mu_p(w,x)\binom{n}{k}}{\binom{|x|+|w|}{k}} = \sum_{i=0}^{n-k-(|J|-|w|)} p^{k+i}(1-p)^{n-k-i}\binom{n}{k+i}\frac{((n-k)-(|J|-|w|))_i}{(n-k)_i}.$$

First note that $\sum_{i=0}^{\log(n)\sqrt{np(1-p)}} p^{k+i}(1-p)^{n-k-i}\binom{n}{k+i}$ tends to $1/2$ by the central limit theorem. Next, for any $i < \log(n)\sqrt{np(1-p)}$ we have $\frac{((n-k)-(|J|-|w|))_i}{(n-k)_i} \sim 1$ and in particular, for sufficiently large n, and i in that range,

$$\frac{((n-k)-(|J|-|w|))_i}{(n-k)_i} > 1/2.$$

So

$$V_w(g) \geq V_w(f)(1/4 - o(1)).$$

□

Using the up-lemma and the down-lemma we now deduce Theorem 7.2 from Theorem 7.1.

Proof of Theorem 7.2. Let p and f be as in the statement of the theorem, let $\varepsilon \geq 0$ and assume n is sufficiently large and $\mathrm{Edge}(f) \leq \delta(\varepsilon)$, where $\delta(\varepsilon)$ is as defined in Theorem 7.1. Let g be as given by the up lemma, Lemma 7.3. Then $\mathrm{Edge}(g) \leq \mathrm{Edge}(f) \leq \delta(\varepsilon)$. Now, invoke Theorem 7.1 to produce $J \subset [n]$ and an intersecting family $T \subset \{0,1\}^J$ which captures g, i.e.,

$$\mathbb{E}_{w \in \{0,1\}^J}[\mathbb{1}_{\overline{T}}(w)\, \mathbb{E}_{x \in \{0,1\}^{[n]\setminus J}}[g(w,x)]] \leq \varepsilon,$$

or, in other words

$$\sum_{w \notin T} V_w(g) \leq \varepsilon\,.$$

By the down lemma, Lemma 7.4, for every $w \in \{0,1\}^J$ (and specifically for $w \notin T$) we have

$$V_w(f) \leq 5V_w(g)$$

so

$$\sum_{w \notin T} V_w(f) \leq 5\varepsilon$$

as required. □

A Theorem 2.1

Theorem 2.1 is basically the main result of [DFR08], apart from some minor differences, the most significant of which being that we improve the quantitative dependence of the parameters (namely, the functions δ_1 and j_1) using the work of Dinur and Shinkar [DS10]. For the reader's convenience, we include a proof sketch in Section A.1.

Alternatively, we now explain how to derive Theorem 2.1 from the original statement in [DFR08], which now follows.

Theorem A.1 ([DFR08, Theorem 1.1 + 2nd and 4th remarks there])**.** *For all $\varepsilon > 0$ there exist $\delta > 0, j \geq 1$, such that the following holds. For all $n \geq 1$ and $f : V^n \to \{0,1\}$ such that $\langle f, Af\rangle \leq \delta$, there exist $J \subseteq [n]$, $|J| \leq j$, and $T \subseteq V^J$ such that*

$$\mathbb{E}_{x \in V^J}[\mathbb{1}_{\overline{T}}(x)\, \mathbb{E}[f(x,\cdot)]] \leq \varepsilon,$$

and $\langle \mathbb{1}_T, A\mathbb{1}_T\rangle \leq \varepsilon$.

The differences between this and our Theorem 2.1 are as follows. First, our theorem considers functions with range $[0,1]$ as opposed to $\{0,1\}$. The proof in [DFR08] actually applies to the more general case, as is easy to check. Alternatively, one can derive the more general case from the restricted one by replacing a function $f : V^n \to [0,1]$ with the function $f' : V^{n+m} \to \{0,1\}$ where we define $f'(x,y)$ to be 1 with probability $f(x)$ and 0 otherwise, independently over all x, y. Then as m goes to infinity, $\langle f', Af' \rangle$ converges to $\langle f, Af \rangle$ and similarly for the other expressions appearing in the theorem.

A second difference is that our theorem involves two functions f_1, f_2 as opposed to just one as above. The proof in [DFR08] can easily be modified to handle this. Alternatively, as before, we can derive this from the original statement as follows. Let a_1, a_2 be two elements of V^2 that are connected by an edge and have no self loops. (Such two elements must exist unless we are in the case in which there is a loop on all vertices in V, which means there are no non-empty independent sets in any power of V, so this case is irrelevant for our current discussion.) Then given $f_1, f_2 : V^n \to [0,1]$ we define $f : V^{n+2} \to [0,1]$ by taking $f(a,x)$ to be $f_1(x)$ if $a = a_1$, $f_2(x)$ if $a = a_2$, and 0 otherwise. Then $\langle f, Af \rangle = 2w\langle f_1, Af_2 \rangle$ where w is the weight of the edge connecting a_1 to a_2.

The final and most significant difference is that Theorem A.1 does not explicitly specify the dependence of δ and j on ε. Inspecting the proof in [DFR08] reveals that the dependence is super-polynomial. By using an improvement by Dinur and Shinkar [DS10] of the technical statement from [DMR09], we are able to obtain a polynomial dependence of the parameters, as stated in Theorem 2.1. We remark that this improvement has no noticeable effect on the final bound in our main result and we could have used the original bound implicit in [DFR08]; we decided to include the improvement as it might be useful for future work.

In slightly more detail, the parameters in the proof of Theorem 1.1 in [DFR08] all depend polynomially on the functions τ_{MOO} and δ_{MOO} defined in Theorem 2.2 there. Those functions can be taken to be polynomial, as shown in the following lemma.

Lemma A.2 (Quantitative version of [DFR08, Theorem 2.2])**.** *There exist functions $\delta_{\mathrm{MOO}}(\varepsilon) > 0$ and $\tau_{\mathrm{MOO}}(\varepsilon) > 0$ such that for any $\varepsilon > 0$, $n \geq 1$, and functions $g_1, g_2 : V^n \to [0,1]$ with $\mathbb{E}[g_1] \geq \varepsilon$, $\mathbb{E}[g_2] \geq \varepsilon$ and $\langle g_1, Ag_2 \rangle < \delta_{\mathrm{MOO}}(\varepsilon)$, there exists a coordinate i with influence greater than $\tau_{\mathrm{MOO}}(\varepsilon)$ on both functions, i.e.,*

$$\mathrm{Inf}_i(g_1) > \tau_{\mathrm{MOO}}(\varepsilon) \ \text{ and } \ \mathrm{Inf}_i(g_2) > \tau_{\mathrm{MOO}}(\varepsilon).$$

Moreover, one can take $\delta_{\mathrm{MOO}}(\varepsilon) = \varepsilon^c$ and $\tau_{\mathrm{MOO}}(\varepsilon) = \varepsilon^c$ for some constant $c > 0$ depending only on A.

This lemma is an immediate corollary of [DS10, Theorem 3.1], and is derived in precisely the same way that [DFR08, Theorem 2.2] is derived from [DMR09, Theorem 3.1]. (In fact, a statement nearly identical to Lemma A.2 already appears as Corollary 3.2 in [DS10].) In more detail and using the notation from [DS10], to derive the lemma above, apply [DS10, Theorem 3.1] in the contrapositive with some $\rho' > \rho$, say $\rho' = \rho^{1/2}$, notice that $\langle F_\varepsilon, U_{\rho'}(1 - F_{1-\varepsilon}) \rangle_\gamma$ is at least ε^C by [DS10, Eq. (2)], and that by definition $\mathrm{Inf}_i^{\leq k}(f) \leq \mathrm{Inf}_i(f)$.

A.1 Proof of Theorem 2.1

Here we include a proof sketch of Theorem 2.1, closely following the original proof in [DFR08] and occasionally borrowing from the notation there.

Claim A.3 ([DFR08, Lemma 2.3.6]). *For any $\eta \in [0,1]$ and function $f : V^n \to [-1,1]$, the sum of influences of the "noisy function" $N_\eta f$ satisfies*

$$\sum_{i=1}^{n} \mathrm{Inf}_i(N_\eta f) \leq (1-\eta^2)^{-2} .$$

In particular, the number of variables that have influence at least τ on $N_\eta f$ is at most $(1-\eta^2)^{-2}/\tau$.

Lemma A.4 (Two-function variant of [DFR08, Lemma 2.5]). *Let $\lambda = \lambda(A) < 1$ be the second absolute eigenvalue of A, and let $1-\lambda < \eta < 1$ be sufficiently close to 1 so that*

$$(1-\eta)\log_\lambda(1-\eta) \leq \sqrt{1-\eta} . \tag{11}$$

Then for any $f_1, f_2 : V^n \to [-1,1]$,

$$\left|\langle f_1, Af_2\rangle - \langle g_1, Ag_2\rangle\right| \leq \sqrt{1-\eta} ,$$

where $g_i = N_\eta f_i$.

Proof. By decomposing the functions according to the eigenbasis of A,

$$\begin{aligned}
\left|\langle f_1, Af_2\rangle - \langle g_1, Ag_2\rangle\right| &= \left|\sum_S (\hat{f}_1(S)\hat{f}_2(S) - \hat{g}_1(S)\hat{g}_2(S))\lambda_S\right| \\
&= \left|\sum_S \hat{f}_1(S)\hat{f}_2(S)(1-\eta^{2|S|})\lambda_S\right| \\
&\leq \sum_S |\hat{f}_1(S)\hat{f}_2(S)| \cdot \max_S (1-\eta^{2|S|})|\lambda_S| .
\end{aligned}$$

A straightforward calculation (see [DFR08, Lemma 2.5]) shows that the above maximum is at most $\sqrt{1-\eta}$. We can therefore complete the proof by noting using Cauchy-Schwarz that $\sum_S |\hat{f}_1(S)\hat{f}_2(S)| \leq \|f_1\|_2\|f_2\|_2 \leq 1$. □

Lemma A.5 ([DFR08, Lemma 2.8]). *There exists a $p = p(A) > 2$ such that the following holds. For any $\varepsilon > 0$, $j, \ell \geq 1$, and any function L mapping each vertex $a \in V^j$ to a subset $L(a) \subseteq \mathbb{N}$ with $|L(a)| \leq \ell$ and satisfying that for at least an ε measure of pairs (a,b) in V^j, $L(a) \cap L(b) \neq \emptyset$, there exists an $i \in \mathbb{N}$ such that*

$$\mu(\{a \in V^j : i \in L(a)\}) \geq (\varepsilon/\ell^2)^{2p/(p-2)} .$$

Lemma A.6 ([DFR08, Claim 3.1]). *For any $\varepsilon > 0$, $\eta \leq 1$, $j \geq 1$, and $f : V^j \to [0,1]$,*

$$\mathbb{E}_{x \in V^j}[\mathbb{1}_{N_\eta f(x) \leq \varepsilon} f(x)] \leq \varepsilon .$$

Proof of Theorem 2.1. Let $p = p(A)$ be as in Lemma A.5 and $\lambda = \lambda(A)$ be the second absolute eigenvalue of A. Fix some $\varepsilon > 0$, and let $\tau_{\mathrm{MOO}}(\varepsilon)$ and $\delta_{\mathrm{MOO}}(\varepsilon)$ be as given in Lemma A.2. Choose $\eta < 1$ close enough to 1 so that $\eta > 1-\lambda$, $2\sqrt{1-\eta} \leq \delta_{\mathrm{MOO}}(\varepsilon)\varepsilon/2$, and Eq. (11) holds. Let $c > 0$ be large enough so that $\delta := \varepsilon^c < \sqrt{1-\eta}$. Define

$$\ell = \frac{2(1-\eta^2)^{-2}}{\tau_{\mathrm{MOO}}(\varepsilon)} ,$$

and choose $\gamma > 0$ small enough so that

$$2\gamma < \tau_{\mathrm{MOO}}(\varepsilon) \cdot (\varepsilon/2\ell^2)^{2p/(p-2)} .$$

For $i \in \{1,2\}$, define $g_i = N_\eta f_i$. Let j be the number of variables with influence greater than γ on either g_1 or g_2, and assume without loss of generality that these are the variables $J := \{1,\ldots,j\}$. By Claim A.3 there are at most $2(1-\eta^2)^{-2}/\gamma$ such variables, so in particular, we can take $j = \varepsilon^{-c}$ for large enough c, as required.

For $a \in V^j$ define $g_{1,a} : V^{n-j} \to [0,1]$ by $g_{1,a}(x) = g_1(a,x)$ and similarly for g_2. Let

$$T_1 = \{a : \mathbb{E}_x[g_{1,a}(x)] \geq \varepsilon\} \subseteq V^j ,$$

and similarly define T_2 with $g_{2,a}$. The condition in Eq. (1) now follows from Lemma A.6.

It remains to prove that $\langle \mathbb{1}_{T_1}, A\mathbb{1}_{T_2}\rangle \leq \varepsilon$. Assume towards contradiction that $\langle \mathbb{1}_{T_1}, A\mathbb{1}_{T_2}\rangle > \varepsilon$. Equivalently, the measure of pairs (a,b) such that $a \in T_1$ and $b \in T_2$ is greater than ε. Notice that by Lemma A.4,

$$\langle g_1, Ag_2\rangle \leq \langle f_1, Af_2\rangle + \sqrt{1-\eta} \leq \delta + \sqrt{1-\eta} \leq 2\sqrt{1-\eta} .$$

Therefore, the measure of pairs (a,b) for which $\langle g_{1,a}, Ag_{2,b}\rangle \geq \delta_{\mathrm{MOO}}(\varepsilon)$ is at most $2\sqrt{1-\eta}/\delta_{\mathrm{MOO}}(\varepsilon) \leq \varepsilon/2$. It follows that there is at least an $\varepsilon/2$ measure of pairs (a,b) for which $\mathbb{E}_x[g_{1,a}(x)] \geq \varepsilon$, $\mathbb{E}_x[g_{2,b}(x)] \geq \varepsilon$, and $\langle g_{1,a}, Ag_{2,b}\rangle < \delta_{\mathrm{MOO}}(\varepsilon)$. By Lemma A.2, for each such pair (a,b), there exists an $i \in \{j+1,\ldots,n\}$ whose influence on both $g_{1,a}$ and $g_{2,b}$ is greater than $\tau_{\mathrm{MOO}}(\varepsilon)$. We can now apply Lemma A.5 with the sets

$$L(a) = \{j < i \leq n : \max(\mathrm{Inf}_i(g_{1,a}), \mathrm{Inf}_i(g_{2,a})) > \tau_{\mathrm{MOO}}(\varepsilon)\} ,$$

whose cardinality is at most ℓ by Claim A.3, and obtain that there exists an $i \in \{j+1,\ldots,n\}$ for which

$$\mu(\{a \in V^j : i \in L(a)\}) \geq (\varepsilon/2\ell^2)^{2p/(p-2)} > 2\gamma/\tau_{\mathrm{MOO}}(\varepsilon) .$$

From this it follows that i has influence greater than γ on either g_1 or on g_2, in contradiction to the definition of J. □

Acknowledgements:

We thank Michael Overton for developing HANSO and Pooya Hatami for sending us an early draft of their work [HST16].

References

[DF09] I. Dinur and E. Friedgut. Intersecting families are almost contained in juntas. *Combinatorics Probability and Computing*, 18(1–2):107–122, 2009.

[DFR08] I. Dinur, E. Friedgut, and O. Regev. Independent sets in graph powers are almost contained in juntas. *GAFA*, 18(1):77–97, 2008.

[DMR09] I. Dinur, E. Mossel, and O. Regev. Conditional hardness for approximate coloring. *SIAM J. Comput.*, 39(3):843–873, 2009.

[DS10] I. Dinur and I. Shinkar. On the conditional hardness of coloring a 4-colorable graph with super-constant number of colors. In *Approximation, randomization, and combinatorial optimization*, volume 6302 of *Lecture Notes in Comput. Sci.*, pages 138–151. Springer, Berlin, 2010. doi:10.1007/978-3-642-15369-3_11.

[DT16] S. Das and T. Tran. Removal and stability for Erdős-Ko-Rado. *SIAM J. Discrete Math.*, 30(2):1102–1114, 2016.

[Fox11] J. Fox. A new proof of the graph removal lemma. *Ann. of Math. (2)*, 174(1):561–579, 2011.

[HST16] P. Hatami, S. Sachdeva, and M. Tulsiani. An arithmetic analogue of Fox's triangle removal argument. *Online J. Anal. Comb.*, (11):Art. 1, 17, 2016.

[KSV12] D. Král', O. Serra, and L. Vena. A removal lemma for systems of linear equations over finite fields. *Israel J. Math.*, 187:193–207, 2012.

[KSV13] D. Král', O. Serra, and L. Vena. On the removal lemma for linear systems over abelian groups. *European J. Combin.*, 34(2):248–259, 2013.

[MOO10] E. Mossel, R. O'Donnell, and K. Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. *Ann. of Math. (2)*, 171(1):295–341, 2010.

[RS78] I. Z. Ruzsa and E. Szemerédi. Triple systems with no six points carrying three triangles. In *Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976), Vol. II*, volume 18 of *Colloq. Math. Soc. János Bolyai*, pages 939–945. North-Holland, Amsterdam-New York, 1978.